From: E. Kwerel,
To: P. Milgrom
Date: June 16, 2000

Questions about Paul Milgrom's Letter of June 14, 2000
re: Package Bidding Public Notice

Benchmark Proposal

1. Is it scalable? Would economically important strategies be ruled out if there were many licenses?
2. What are the reasons for the limitation that no bidder may bid on more than 24 combinations of licenses during the full course of the auction? Is this both for computational simplicity and to prevent adverse strategic behavior such as excessive parking? Could you give a simple example of how increasing the number of allowed combinations increases the parking problem or opportunity for other adverse strategic behavior.
3. The formula for calculating the minimum acceptable bid on each package includes the highest previous bid on the package, plus a percentage increment. How is the "highest previous bid" defined?
   a) Is it based on the best combination of non-overlapping bids on that package considering all bids on that package *and* subsets of the package?
   b) Suppose that the provisionally winning bid is less than the previous high bid (because the previous high bid is a mutually exclusive bid with some other bid that is provisionally winning). Can a bidder submit a new bid that is higher the provisionally winning bid but lower than the previous high bid?
4. Could we substitute the "shortfall" rule for calculating the minimum accepted bid (proposed in the "Refining the Public Notice" section of the letter) for the rule (in the "Benchmark" section) that includes using the RAD approach?
5. If we use the "shortfall" rule how should it be calculated? Was McAfee objecting to the percentage of the "deficit" that would need to be made up (he suggested 30% instead of 50%) or was he also suggesting that the percentage be applied to the "shortfall" instead of the "deficit"? You proposed that it be applied to the deficit, which is the shortfall multiplied by the proportion of bidding units associated with the combination. Exactly how is the deficit calculated? For a particular package that is not part of the provisionally winning set do we calculate the maximum revenue that could be generated with that package and then allocate the shortfall to that package as a *percentage of* the bidding units for all the *non-provisionally winning bids*?
6. Is the activity rule stringent enough? Should we consider the 50% activity rule? Would a rule that requires a bidder to maintain activity on some minimum percentage of bidding units for which it is eligible force bidders to start on larger packages and move to smaller packages as a back up strategy? Is this a problem?

Refining the Public Notice Design

7. Are all bids "kept" under this proposal in the sense that they are available for inclusion in the provisionally winning set? Do some of the kept bids not get activity credit?
8. How do we determine activity credit in the general case with more than the nine packages?
9. Clarify what a bidder must do to "renew" a retained bid in the current round. Is it entered at the price at which it was retained? In the letter it says that it is entered "at the same price as its current provisionally winning bid." What does this mean?
10. You state that the sole cogent argument against prohibiting cancellation of non-winning bids is that it prevents straightforward bidding. What are the other arguments for allowing bid cancellation? Plott seems to argue that allowing cancellation lets bidders explore possibilities and actively helps them better discover winning coalitions. The Weber statement seems to embody a similar view.

Following Response from P. Milgrom to E. Kwerel on June 17, 2000

Paul Milgrom's Answers to Questions about His Letter of June 14, 2000
re: Package Bidding Public Notice

Benchmark Proposal

1. Is the benchmark proposal scalable?

I am not certain if it is "scalable." The following discussion lays out the scalability issues as they presently appear to me.

From a computational perspective, the problem of finding an optimum in a combinatorial auction with unlimited packages is hard—"NP-complete" in the classification used by computer scientists. This means, roughly, that the time needed to compute a solution can grow exponentially with the number of bits of data required to describe the problem. This formal conclusion has been applied less formally to conclude that when the number of licenses is large, the problem becomes at least potentially very hard. That, in turn, is what analysts usually mean when they say that general combinatorial auctions are not "scalable."

This negative conclusion, however, is not completely convincing to me. First, the presence of a large number of licenses does not imply that the problem takes many bits to describe if the number of relevant license combinations remains small. That might happen even in an auction with many licenses if evaluating packages requires formulating business plans for the package, because doing that is costly. To put it another way, the auction designer can guarantee that the problem stays manageable by limiting bidders to bid on a small number of packages of their own choosing, as in the benchmark proposal.

The second factor that casts doubt on scalability concerns is the inexorable growth of computing capacity. The balance between the costs of computing provisional solutions in the auction and the costs of estimating values is shifting rapidly. According to "Moore's Law," computing speeds double every eighteen months, making computability issues much less significant as time passes. Improved combinatorial optimization algorithms are contributing further to this shift in the balance.

The most important scalability advances in the future will be in finding ways to allow bidders to express bids compactly and to track intuitively what is going on during the auction. It is the user experience, rather than computation, that is rapidly becoming the most important limitation in scaling combinatorial auctions.

2. Would economically important strategies be ruled out if there were many licenses and bidders were limited to bidding on 24 license combinations?

Based on my own experience consulting with bidders, I believe that allowing 24 combinations will not rule out any important strategies for the 700 MHz auction, because real bidders will not consider that many alternative business plans.

While the general approach makes sense, there are auction situations in which bidders would consider more than 24 alternatives. Imagine, for example, a situation in which there are two types of licenses, A and B, and some bidders

want one "type A" license and one "type B" license. Suppose there are six licenses of each type, each with slightly different characteristics. In that case, a bidder might evaluate all 6x6=36 alternatives that involve combinations of A and B.

3. What are the reasons for the limitation that no bidder may bid on more than 24 combinations of licenses during the full course of the auction? Is this both for computational simplicity and to prevent adverse strategic behavior such as excessive parking? Please give a simple example of how increasing the number of allowed combinations increases the parking problem or opportunity for other adverse strategic behavior.

Limiting bidders to 24 combinations does have real advantages for easing computations, but in the 700 MHz auction, these computational concerns are secondary: the main reason for imposing this limit is to limit parking opportunities and to make the auction easier for all bidders to follow. As previously noted, there are 4095 combinations of twelve licenses, and it is likely that most of those combinations will attract few serious bids in the auction. Bidders who wish to delay making serious bids may find lots of opportunities among these packages.

I decline to provide a specific example, because the available parking strategies will depend on the other rules, including especially the rules governing activity and minimum bids. It would be misleading to give an example that is "corrected" by activity or minimum bid rules in the particular case. As a general matter, allowing additional combinations can only increase the number of "parking" bids available to bidders.

4. The formula for calculating the minimum acceptable bid on each package includes the highest previous bid on the package, plus a percentage increment. How is the "highest previous bid" defined?

    a) Is it based on the best combination of non-overlapping bids on that package considering all bids on that package *and* subsets of the package?
    b) Suppose that the provisionally winning bid is less than the previous high bid (because the previous high bid is a mutually exclusive bid with some other bid that is provisionally winning). Can a bidder submit a new bid that is higher the provisionally winning bid but lower than the previous high bid?

(a) No. This rule would exclude bids that can eventually become winners of the auction.

(b) This proposal is different from what I had suggested for the benchmark, but is superior to it. The theoretically best rule makes each bidder's minimum bid for a package in the round equal to the larger of a threshold (such as the one discussed in the next question) or that bidder's own previous best bid for that package plus a percentage.

5. Could we substitute the "shortfall" rule for calculating the minimum accepted bid (proposed in the "Refining the Public Notice" section of the letter) for the rule (in the "Benchmark" section) that includes using the RAD approach?

Yes.

6. If we use the "shortfall" rule how should it be calculated?
    (a) Was McAfee objecting to the percentage of the "deficit" that would need to be made up (he suggested 30% instead of 50%) or was he also suggesting that the percentage be applied to the "shortfall" instead of the "deficit"? You proposed that it be applied to the deficit, which is the shortfall multiplied by the proportion of bidding units associated with the combination.
    (b) It is unclear exactly how the deficit should be calculated. For a particular package that is not part of the provisionally winning set, do we calculate the maximum revenue that could be generated with that package and then allocate the shortfall to that package as a *percentage of* the bidding units for all the bids that are NOT *provisionally winning bids*?

    (a) The definitions of the terms "deficit" and "shortfall" were added to the letter after McAfee's comments were received, in an attempt to add precision to the discussion. I cannot be certain which of these concepts McAfee had in mind when he made his remarks.
    (b) Yes.

7. Is the activity rule stringent enough? Should we consider the 50% activity rule? Would a rule that requires a bidder to maintain activity on some minimum percentage of bidding units for which it is eligible force bidders to start on larger packages and move to smaller packages as a back up strategy? Is this a problem?

    Activity rules are an area of the benchmark that remains underdeveloped due to the limited time to consider them thoroughly. These rules and minimum bid rules need to be considered together to determine their adequacy.

    In the benchmark auction, an important new point arises. Because bids are mutually exclusive in the benchmark auction, alternative strategies can safely be pursued in parallel. In contrast, in the simultaneous multiple round auction, alternative strategies have to be pursued in sequence. Moreover, in the benchmark auction, straightforward bidders do not skip rounds in bidding on any particular package (except when the package is a provisional winner). These facts suggest new kinds of activity rules, such as one that prevents a bidder from being intermittently active on any package. Activity rules based on this approach would seem to allow bidders ample flexibility to bid in approximately straightforward ways. Combined with suitably aggressive minimum bids, they would also assure a quick pace for the auction.

    Assuming that the 700 MHz auction will need to use a variant of the established activity rules, the proposed 50% activity rule based on points adds a very important margin of security to auction design. If the points represent a reasonable estimate of value ratios, then the threat to efficiency by imposing such a rule are likely to be small.

Refining the Public Notice Design

8. Are all bids "kept" under this proposal in the sense that they are available for inclusion in the provisionally winning set? Do some of the kept bids not get activity credit?

In view of the insights provided by the benchmark and provided that bids in different rounds by the same bidder are deemed mutually exclusive, "keeping" all bids appears to be a valuable rule modification. Activity credit issues are discussed separately below.

9. How do we determine activity credit in the general case with more than the nine packages?

A prior question is which bids are retained bids. I assume for my answer that the "rows and columns" structure for determining retained bids is preserved.

In that case, based on a suggestion by Professor Charles Plott, I recommend that the auction give 100% credit for provisionally winning bids and new eligible bids and only 50% credit for "retained bids" that are not provisional winners. To ensure timely completion, *the auctioneer should have authority to reduce the activity credit for retained bids that are not provisional winners to 0%.*

10. Please clarify what a bidder must do to "renew" a retained bid in the current round. In the letter, it says that the renewed bid is entered "at the same price as its current provisionally winning bid." What does this mean?

In any round, a bidder may repeat as a current bid any of its retained bids (and any of its bids "kept" in the system). The repeated bid is not subject to the minimum bid requirement and earns no activity credit, but is not regarded as mutually exclusive with other bids at the current round. This rule merely allows a bidder to "waive" the exclusivity associated with its retained bids.

The reference to the "current provisionally winning bid" should be struck.

11. You state that the sole cogent argument against prohibiting cancellation of non-winning bids is that it prevents straightforward bidding. What are the other arguments for allowing bid cancellation? Plott seems to argue that allowing cancellation lets bidders explore possibilities and actively helps them better discover winning coalitions. The Weber statement seems to embody a similar view.

The relevant arguments can be described as follows. A bidder wishes to purchase either package A or package B, but not both, perhaps because it lacks the budgetary resources to buy both. The bidder has a retained bid for package A, but its chance of winning appears to the bidder to be less than 50%. Under the rules in the Public Notice, if the bidder places a bid for B, it faces an exposure problem: it could wind up winning both A and B. This exposure problem would be resolved if a bidder could simply cancel its retained bids.

The cogency of this argument depends on the idea that the bidder has a "legitimate need" to bid for both A and B without facing an exposure problem. We would hardly be concerned about the bidder's predicament if the bid on A were, say, a parking bid. What lends cogency to the argument is that straightforward

bidding could sometimes lead to this situation. I suspect that this is just what Weber and Plott have in mind when they offer their comments.

There may be other patterns of bidding that some commenters consider to justify bid cancellation, but I can find none that I consider convincing. If others disagree, I leave it to them to elaborate the argument.